



US 20020080196A1

(19) **United States**

(12) **Patent Application Publication**  
**BORNSTEIN et al.**

(10) **Pub. No.: US 2002/0080196 A1**

(43) **Pub. Date: Jun. 27, 2002**

(54) **AUTO-SUMMARY OF DOCUMENT  
CONTENT**

(21) Appl. No.: **09/241,255**

(22) Filed: **Feb. 1, 1999**

(76) Inventors: **JEREMY J. BORNSTEIN, MENLO  
PARK, CA (US); DOUGLASS R.  
CUTTING, OAKLAND, CA (US);  
JOHN D. HATTON, MT. HERMON,  
CA (US); DANIEL E. ROSE,  
CUPERTINO, CA (US)**

**Related U.S. Application Data**

(63) Continuation of application No. 08/536,020, filed on  
Sep. 29, 1995, now patented.

**Publication Classification**

(51) **Int. Cl.<sup>7</sup> ..... G06F 3/00**

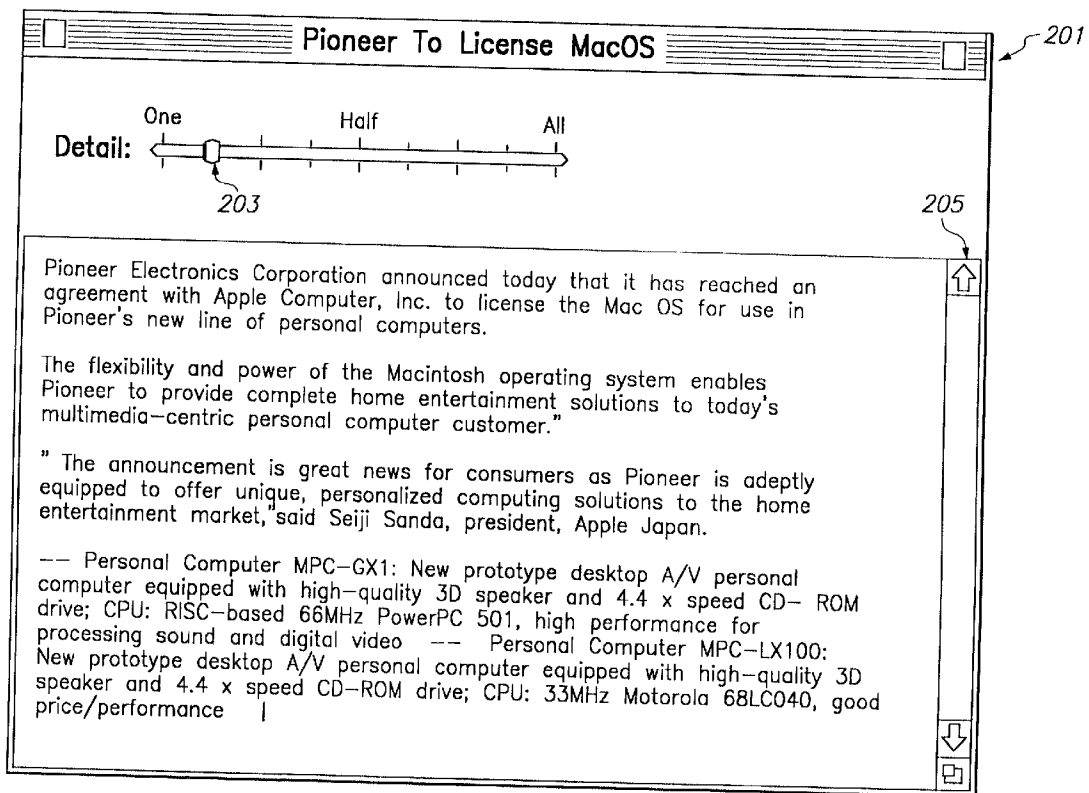
(52) **U.S. Cl. .... 345/854**


Correspondence Address:  
**APPLE COMPUTER INC  
1 INFINITE LOOP  
M/S 38-PAT  
CUPERTINO, CA 95014**

(57) **ABSTRACT**

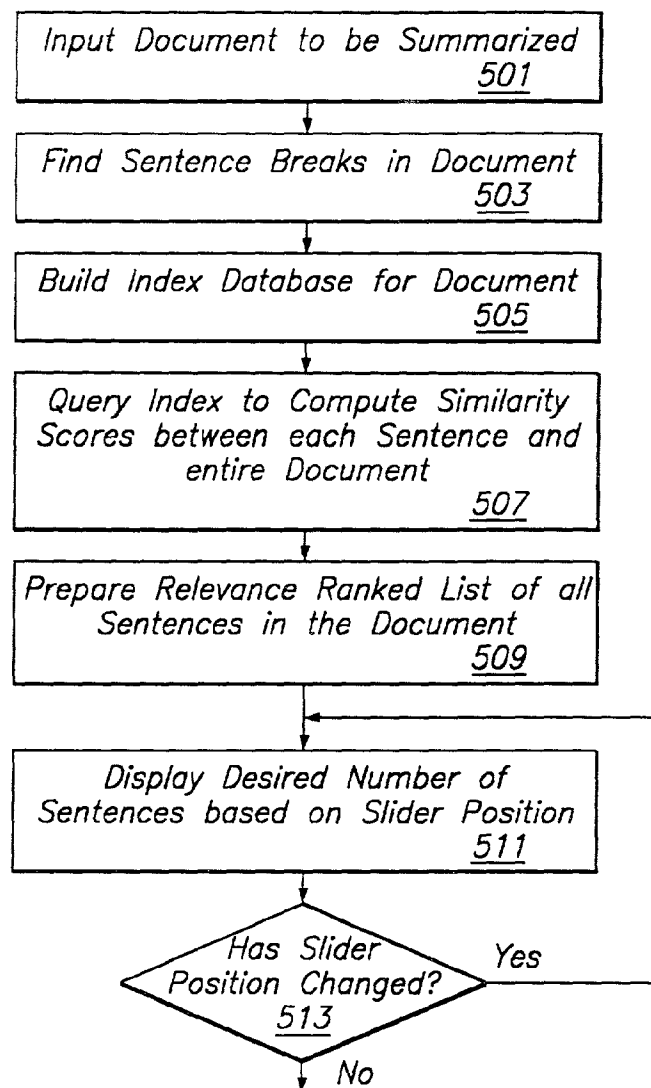
(\*) Notice: This is a publication of a continued prosecution application (CPA) filed under 37 CFR 1.53(d).

A computer system user interface provides a document summary which allows the user to more easily identify the contents and subject matter of the document.





**FIG. 1**



**FIG. 5**

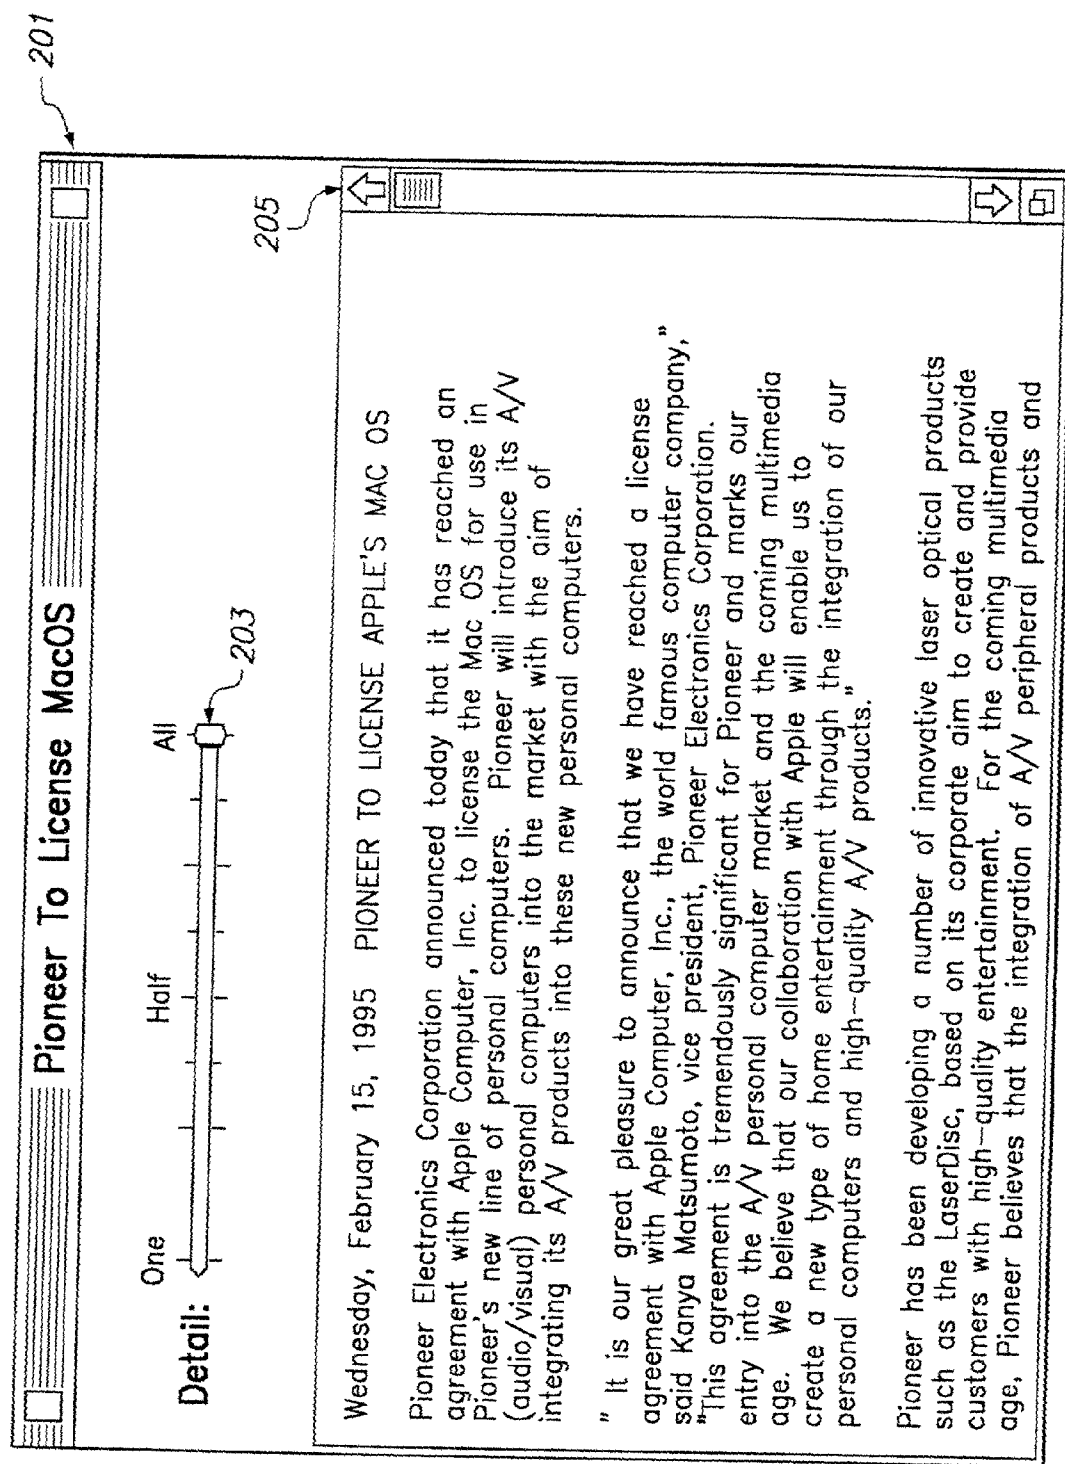


FIG. 2

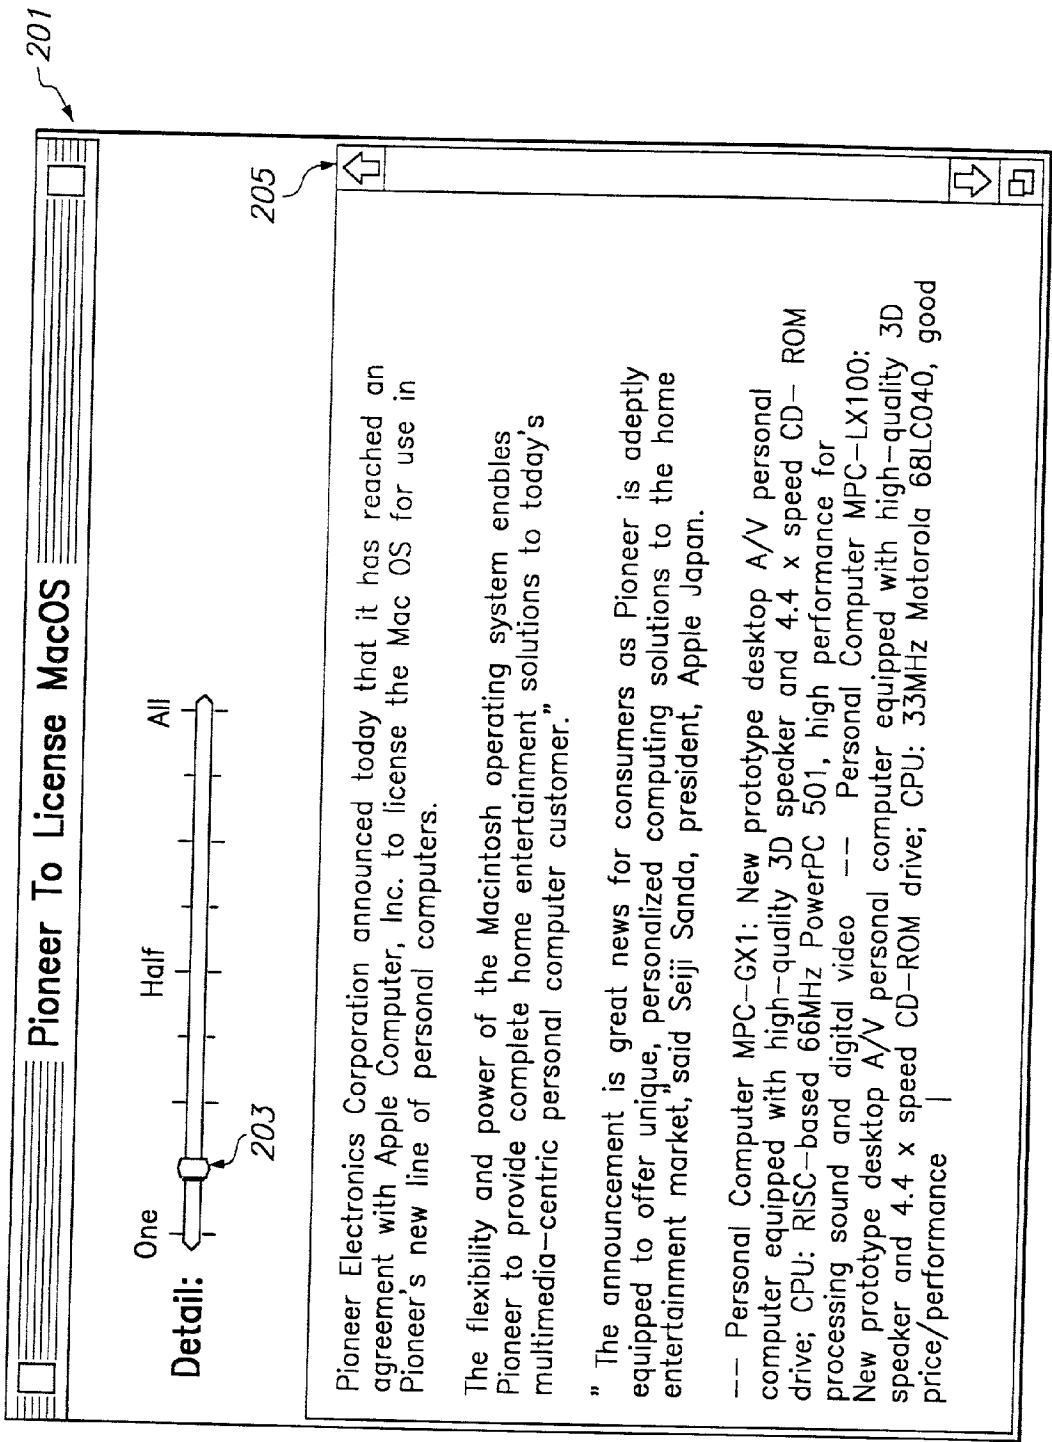


FIG. 3




FIG. 4

Automatic Summary Window				
Name	Size	Kind	Label	Last Modified
Auto Summary Patent	40K	MacWrite	Pro docu...	Tue, Sep 26, 1995, 5:24 PM
Monday Night NFL 9/25/...	16K	MacWrite	Pro docu...	Tue, Sep 26, 1995, 5:11 PM
Pioneer To License MacOS	16K	MacWrite	Pro docu...	Tue, Sep 26, 1995, 4:59 PM
				A real-time interactive documen...
				Detroit over San Francisco 27-2...
				Personal Computer MPC-GX1 - N...

FIG. 6

Auto Summary Patent

MacWrite® Pro Document

Kind: MacWrite Pro document

Size: 40K on disk (29,372 bytes used)

Where: Macintosh HD: Automatic Summary Window:

Created: Tue, Sep 26, 1995, 5:24 PM

Modified: Tue, Sep 26, 1995, 5:24 PM

Version: MacWrite® Pro 1.5v1

Comments: December 1993

A real-time interactive document summarization system which allows the user to continuously control the amount of detail to be included in a document summary.

☐ Locked

☐ Stationery pad

Auto Summary Patent

Monday Night NFL 9/25/...

Pioneer To License MacOS

Comments

A real-time interactive document summarization system which allows the user to continuously control the amount of detail to be included in a document summary.

FIG. 7

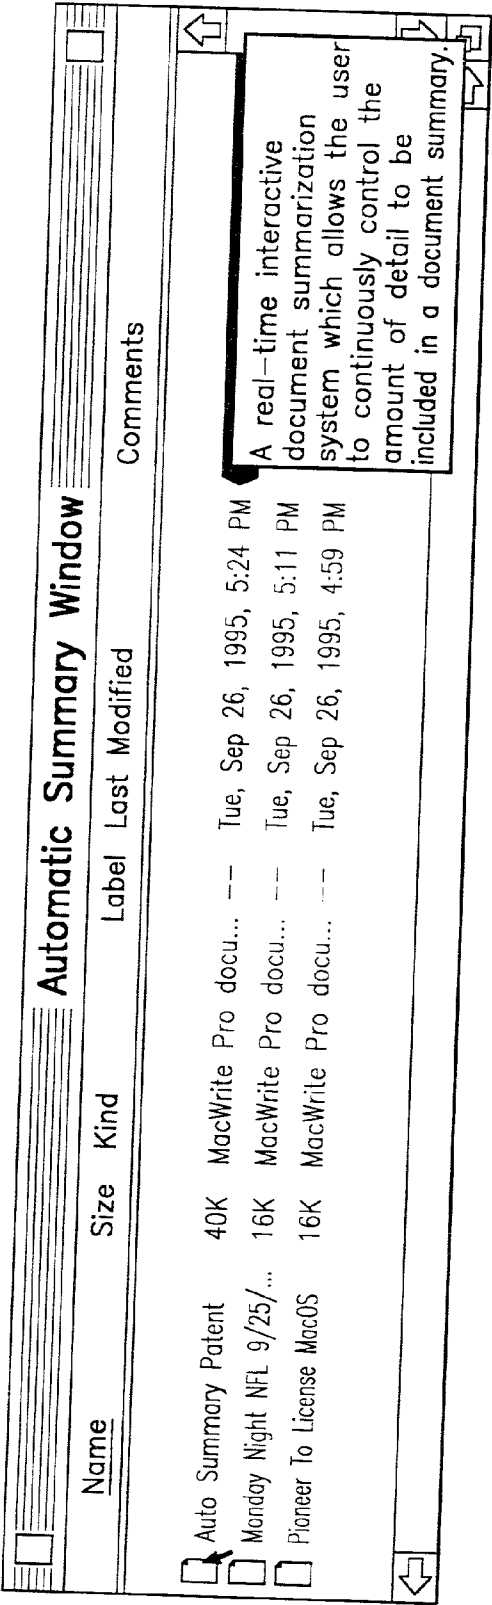


FIG. 8






FIG. 9

## AUTO-SUMMARY OF DOCUMENT CONTENT

### REFERENCES TO RELATED APPLICATIONS/INCORPORATION BY REFERENCE

[0001] This application is a Continuation of application Ser. No. 08/536,020, filed on Sep. 29, 1995, which is incorporated herein by reference. This application is also related to U.S. Pat. No. 5,838,323 which was also filed on Sep. 29, 1995. This patent is also incorporated herein by reference.

[0002] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights.

### FIELD OF THE INVENTION

[0003] The present invention relates to the field of document summarization which is otherwise known as automatic abstracting wherein an extract of a document (i.e., a selection of sentences from the document) can serve as an abstract.

### BACKGROUND OF THE INVENTION

[0004] The advent of the personal computer and modem telecommunications has resulted in millions of computer users communicating with each other around the globe. One of the primary uses of such computers by such users is accessing the vast store of digital information which has been created over the last several decades. Further, additional digital information is created daily due to both the conversion of information previously unavailable digitally and the large amount of new information created by an ever increasing computer user population.

[0005] One concern with this vast, ever increasing amount of digital information is the time it takes to read even a small portion of it. Whether one is reviewing a previously arranged set of documents, as in the case of reading an on-line newspaper or magazine, reviewing the results of an electronic search, or scanning documents stored on a large hard disk drive of a personal computer, it can still take considerable time to read more than a minimal amount.

[0006] What is needed, therefore, is a facility which provides a summary or abstract of each document. Having a summary of each document allows the reader to determine whether that document is of interest, and hence, reading more of the document might be desirable. Conversely, reading the summary of a document could suffice to sufficiently inform the reader about the document, or instead, could indicate to the reader that the particular document is not of interest. No matter the result, a good document abstract mechanism could be quite valuable in the modern digital world.

[0007] However, a good document abstract mechanism means more than merely providing an automatic summary of a document. Prior approaches to document summarization or "Automatic Sentence Extraction", as discussed on pages 87-89 of the "Introduction to Modern Information

Retrieval" by Salton and McGill, Copyright 1983, incorporated herein by reference in its entirety, have yet to yield abstracts "in a readable natural language context" which "obey normal stylistic constraints." Salton and McGill further state that "[r]eadable extracts are obtainable without excessive difficulties, but perfection cannot be expected within the foreseeable future."

[0008] One difficulty with prior document abstract mechanisms, even when overcoming many of the natural language barriers, is that the system or mechanism can never know for certain whether the user is receiving as much or as little of an abstract as they would like. In other words, no matter how well the mechanism can determine which portions of the document to include in the summary or abstract, the mechanism can never automatically include just the right amount of abstract to always please the user. This can be due to different users' interest levels, different user's reasons for reviewing the document, and even time or situation varying interests of the same user. As such, what is needed is not necessarily a better abstracting algorithm as much as a mechanism which allows the user to interactively specify whether the present abstract is sufficient or, instead, whether more or less of the original document should be included in the abstract or summary.

[0009] The present invention utilizes an interactive control which allows the user to specify whether more or less of the original document should be included in the document summary. Allowing the user to interactively control how much of the original document gets included in the summary facilitates rapid review of documents in which the user has little interest as well as review of up to the entire document in the case of great user interest. Furthermore, such interactive control allows the user to expand and contract summarized documents at will, thus freeing the user to focus on the content of the summarized document rather than on trying to determine what amount or percentage is sufficient or how the underlying abstracting mechanism operates.

### BRIEF DESCRIPTION OF THE DRAWINGS

[0010] The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements, and in which:

[0011] FIG. 1 is a diagram of a typical computer system as might be used with the present invention;

[0012] FIG. 2 is a sample summary document window according to one implementation of the present invention wherein "All" of the original document to be summarized is displayable;

[0013] FIG. 3 is a sample summary document window according to one implementation of the present invention wherein one-eighth of the original document to be summarized is displayable;

[0014] FIG. 4 is a sample summary document window according to one implementation of the present invention wherein "One" most representative sentence of the original document to be summarized is displayable;

[0015] FIG. 5 is a flowchart of the document summarization methodology according to one implementation of the present invention;

[0016] FIG. 6 is a sample user interface display showing some or all of the “top sentence” of each document in a display line or listing of documents in a computer system user interface;

[0017] FIG. 7 is a sample user interface display showing the “top sentence” of a document in a comments field of an informational window of the document in a computer system user interface;

[0018] FIG. 8 is a sample user interface display showing the “top sentence” of a document in a pop-up area of a display line or listing of documents in a computer system user interface; and,

[0019] FIG. 9 is a sample user interface display showing the “top sentence” of a document in an open dialog box in a computer system user interface.

#### SUMMARY AND OBJECTS OF THE INVENTION

[0020] It is an object of the present invention to provide an interactive document summarization system.

[0021] It is a further object of the present invention to provide an interactive document summarization system wherein the user of the system can control the amount of the document summary.

[0022] It is a still further object of the present invention to provide a file listing containing document summary information.

[0023] It is an even further object of the present invention to provide document summary information about a document in a variety of contexts.

[0024] The foregoing and other advantages are provided in a computer file system file directory listing displayed on a computer display, wherein a method of displaying a summary of a document comprising one or more sentences comprises, (i) displaying in the computer file system file directory a listing referencing the document, (ii) ranking the relevance of the one or more sentences of the document to the document as a whole, and (i) displaying in the computer file system file directory the highest ranking one or more sentences of the document.

[0025] The foregoing and other advantages are also provided by a document summary display system comprising, (i) a document containing one or more separate sentences, (ii) a relevance ranking means for ranking the relevance of the one or more separate sentences of the document, and, (iii) a display means for displaying a summary of the document within a computer system user interface listing containing a displayed reference to the document wherein the summary of the document is based upon the relevance ranking of the one or more sentences.

[0026] Other objects, features and advantages of the present invention will be apparent from the accompanying drawings and from the detailed description which follows.

#### DETAILED DESCRIPTION OF THE INVENTION

[0027] The present invention can be implemented on all kinds of computer systems. Regardless of the manner in which the present invention is implemented, the basic operation of a computer system embodying the present invention, including the software and electronics which allow it to be performed, can be described with reference to the block diagram of FIG. 1, wherein numeral 10 indicates a central processing unit (CPU) which controls the overall operations of the computer system, numeral 12 indicates a standard display device such as a CRT or LCD, numeral 14 indicates an input device which usually includes both a standard keyboard and a pointer-controlling device such as a mouse, and numeral 16 indicates a memory device which stores programs according to which the CPU 30 carries out various predefined tasks. The interactive document summarization program according to the present invention, for example, is generally also stored in this memory 16 to be referenced by the CPU 10.

[0028] As stated above, the process of document summarization or automatic abstracting is well known in the art. A variety of different mechanisms, used singly and in combination, have been tried to automatically create document summaries or abstracts. Such mechanisms typically start with determining the significance of particular words and/or sentences, usually by focusing on position in the document, semantic relationships, and term frequencies. Further criteria may include contextual inference and/or syntactic coherence.

[0029] However, again, regardless of the sophistication of the summarization mechanism (and note that the present invention is equally applicable to document summarization using any reasonable summarization mechanism now known or later developed), it is highly unlikely that any particular summarization mechanism will always generate the degree of detail desired by the user. As such, the present invention provides the user with a control mechanism to vary the degree of summary detail so as to suit the particular user's tastes and interests at that point in time and for that particular purpose.

[0030] In the preferred embodiment of the present invention a “summarization engine” (again, any reasonable summarization mechanism would work with the present invention) running on a personal computer is used to rank all of the sentences in a document from most to least representative. The user interacts with the system by adjusting a slider control displayed in a graphical user interface of the computer system. As the user moves the slider to a given position, the engine returns the top n sentences, where n is based on the slider's position. The sentences' original order and paragraph structure are maintained in the preferred embodiment of the present invention as a summary consisting of those n sentences is displayed in a window on the computer screen.

[0031] The effect of the present invention is that as the user moves the slider, the window instantaneously updates to display a summary with more or less detail and in the same order as the original document. Thus, as the user moves the slider to ask for more detail the summarized document appears to grow with the ever-increasing number of sentences instantly appearing in their original order and paragraph structure (with the upper limit being the entire original document). And as the user moves the slider to ask for less detail the summarized document appears to shrink with the sentences instantly disappearing and the remaining sentences within each remaining paragraph collapsing to

form new summary paragraphs (with the lower limit being the one sentence most characteristic of the entire document according to the summarization mechanism). And again, the interface mechanism of the preferred embodiment of the present invention operates as simply as having the user manipulate a cursor control device such as a mouse, trackball or trackpad, to move a slider control on the computer display to indicate that more or less summary information is desired.

[0032] Referring now to **FIG. 2**, a sample screen from the system before it has summarized the document can be seen. In the figure, a document summary window **201** can be seen wherein the slider **203** is set to "All," indicating that all of the sentences in the original document are to be shown. The scroll bar **205** on the right hand side of the window, a standard feature of the standard Macintosh Finder user interface environment, indicates that there is more of the document that exists than can fit within the window **201** displayed on the screen (in other words, while the "All" setting allows viewing of the entire document, not all of the document may be displayable at a given point in time due to display screen and/or window size constraints). In this example, the original document contains 32 sentences and, with this window size, would fill several screens of text.

[0033] Referring now to **FIG. 3**, the user has moved the slider **203**, typically via a cursor control device such as a mouse, trackball, or trackpad, to indicate that he only wants a summary one-eighth the size of the original document (note that predetermined summarization settings, wherein the system automatically generates a preset amount of summarization according to previously set system or user values, are equally supportable with the present invention) to be displayed within the document summary window **201**. The summary now fits within the window **201**, as indicated by the empty scroll bar **205** on the right hand side of the summary window.

[0034] Referring now to **FIG. 4**, the user has now moved the slider **203** to indicate that he only wants a summary which shows the one sentence deemed by the summarization engine to be most representative of the document's content to be displayed within the document summary window **201**.

[0035] It is important to note here that the examples of **FIGS. 2-4** are merely static points in time and that the user has the flexibility to continuously alter the slider position. In this way, the user might first see the summary window as it appears in **FIG. 3**, wherein one-eighth of the document is displayed. Then, the user might continuously move the slider towards the "All" setting thus requesting more and more of the document be displayed in the summary window until he reaches the summary window as it appears in **FIG. 2**, wherein all of the original document is available for viewing. Then, the user might decide that less of the document is desired to be viewed and thus move the slider back towards the "One" setting, such that the system is continuously showing less and less of the original document. Finally, the user might end up moving the slider all the way down to the "One" setting, wherein only the one most indicative sentence is displayed in the document summary window as it appears in **FIG. 4**.

[0036] As just explained, a significant advantage of the present invention lies in the use of the slider or knob user interface control. Just as in the case of a dimmer switch to

control room lighting, which provides direct-feedback by having the light get brighter or dimmer as the user moves the slider or knob control as well as having an essentially infinite number of settings, using a slider or knob control in the present invention has greater intuitiveness and utility than would mere up and down buttons having discrete, quantized levels. A slider control combined with immediate display feedback (immediately displaying greater or fewer sentences in the document being summarized as the user moves the slider) means the user only has to be concerned about whether the amount of summarized information being displayed is of the desired quantity.

[0037] And the present invention has clear advantages over requiring the user to specify actual summary values or percentages. Just as in the case of a light dimmer switch where the user only knows that they want more or less light (rather than, say, knowing that what they want is 15% more light or 22% less light), the slider control of the present invention avoids placing on the user the additional cognitive load of first estimating the new amount desired. In other words, after the user determined that more or less summary information was desired, if the interface mechanism required specifying a summary percentage or utilizing up and down buttons then the user would have to be concerned with exactly how much or less information is truly desired. It is less intuitive to require the user desiring more information to first determine that 49% isn't enough but that 58% is sufficient or to try a series of static up and down clicks until the desired amount is obtained. The more intuitive interaction mechanism of the present invention allows the user to interactively operate a continuously variable control while providing immediate display feedback of the greater or lesser information until the user determines that the appropriate amount of information is displayed.

[0038] Thus, another advantage of the present invention, as alluded to above, is that the user has the option of continuously changing the amount of summary information being displayed which thus facilitates the user requesting more and more of the original document as the greater and greater summary amount further piques the user's interest. And then, after the user has read the desired amount of document summary, the user still has the option of decreasing the final amount of summary information. This has the added benefit of providing the reader with as much information as desired while still facilitating minimal document summaries which might then be used in other ways (e.g., see below regarding "View by Sentence" and "comment window" applications).

[0039] A general overview of the summarization engine of the present invention will now be explained. Note first, however, that any of a large variety of well-known summarization techniques are equally applicable to the present invention. In many prior art document retrieval systems a "vector model" approach has been taken where each record or document is represented by a vector representative of the distribution of terms in the document. A particular search query is then represented as a vector such that the retrieval of a particular record or document then depends upon the magnitude of a similarity computation between the particular document's representative vector and the query's representative vector. Suffice it to say that the vector model of document comparison is well known in the art of computer search and retrieval mechanisms (see Salton and McGill.

Introduction to Modern Information Retrieval. 1983, pages 120-123, Salton and Buckley, "Term-Weighting Approaches in Automatic Text Retrieval" Information Processing & Management, Vol. 24, No. 5, pp. 513-523. Witten, Moffat, and Bell. Managing Gigabytes, Compressing and Indexing Documents and Images, 1994, pp. 141-148, and Frakes and Baeza-Yates, Information Retrieval, Data Structures & Algorithms, 1992, pp. 363-392, all incorporated herein by reference in their entirety).

[0040] Typical prior art search and retrieval mechanisms, however, attempt to find, out of a corpus comprised of multiple documents, one or more documents which are most similar to a single query which may itself be a document. Instead, the preferred embodiment of the present invention treats each sentence in the document to be summarized as being equivalent to an entire document, and thus the set of all of the sentences of the document can be treated as the corpus of documents to be searched. Then, the present invention treats the text of the original document as the query to be applied to the corpus. In this way, a determination can be made as to how similar each sentence in the document is to the document as a whole. The result is a ranking or value score for each sentence in the document being summarized. Then, depending upon either a preset value *n* or the user specified slider setting *n*, only those sentences above the ranking or value score of *n* get displayed in the document summary.

[0041] Furthermore, the present invention, as is common in the art, uses term weighting to provide distinctions between the various terms or, in the present invention, words in a document. The present invention utilizes a well known term weighting formula (see, e.g., page 518 of Salton and Buckley in the "Term-Weighting Approaches in Automatic Text Retrieval" article referred to above and incorporated herein) wherein the term-weighting components are as follows:

[0042]  $tf$  = the number of times a term (word) occurs in a sentence or in a document as a whole;

[0043]  $N$  = the number of sentences in the document; and,

[0044]  $n$  = the number of sentences in the document which contain a given term.

[0045] The term-weighting formula is applied to both document and query vector terms and is  $tf$  where  $t$  is replaced by  $\log(tf+1)$  to better normalize long documents and to keep things positive,  $f$  is replaced with  $\log(N/n)+1$  to permit a search for a word that occurs in every sentence to in fact find every sentence, and  $c$  is unaltered, i.e., each weight in a vector is divided by the square root of the sum of all the squares of the unnormalized weights for the vector.

[0046] Referring now to FIG. 5, the process of the present invention will now be described. When a document is to be summarized 501 with the present invention, it must first be determined 503 where the sentence breaks are in the document. Note that the sentence break determination approach of the preferred embodiment of the present invention is shown in the C++ programming language format in Appendix A to the present specification.

[0047] The next step is to determine the sentence ranking within the document being summarized. This is accomplished by first 505 building an index which is a database representing the contents of the sentences in the document in the form of statistics about the words in those sentences, a process which is well known in the art. Then, 507 the entire original document is treated as a query to the corpus of individual sentences in the document in accordance with the standard vector model approach. The result is a score indicating how well each sentence matches the query of the entire document and, hence, the output of the queries is a rank ordered list by score of all the sentences in the document 509.

[0048] Then, the desired number of sentences to include in the document summary display is determined 511, once a ranked list of each sentence in the original document is obtained, by examining either a preset value or the slider position value which thus indicates how far down the ranked list to go. Again, the markers on the slider could be represented as a proportional amount of the entire document, as a numeric value of the number of sentences of the total document, or even as a non-linear value indicator of the total document. While this last form may not sound as intuitive as the former ones, it is important to note that studies have shown that most of the content of a document can be understood by only reading a relatively small amount of the entire document (e.g., 20-25%). Further, remember that the user interface of the present invention frees the user to focus on the displayed summary content rather than on some more obscure summary percentage or value. As such, a non-linear slider may provide even greater utility to the user of the present invention.

[0049] Lastly, the slider position is monitored 513 so that if the user changes its position, thus indicating a desire for more or less information, the appropriate amount of summary information based on the new slider position 511 can be displayed.

[0050] It is important to note a performance advantage in the process just described. In the preferred embodiment of the present invention, because the query 507 asked for all of the sentences in the document before concerning itself with how many sentences will be displayed, every sentence in the document gets a ranking 509. Then, whenever the slider position is changed 513, displaying the larger or smaller summary is a relatively simple matter of merely displaying the more or less sentences as dictated by the previously generated relevance ranked list. In other words, by precomputing the relevance ranking, displaying more or less detail can be accomplished quickly without an additional query to be performed for each change in the slider position.

[0051] Further, in the preferred embodiment of the present invention, displaying more or less detail is done using an offscreen bitmap, a technique well known in the computer art. Using an offscreen bitmap makes the display appear to have the sentences instantly inserted or deleted in place rather than having the entire document summary appear to scroll from the top down whenever the user asks for more or less detail.

[0052] Note that the present invention has numerous applications. A more clear application would be as part of a document browser or within a document retrieval context thus allowing more rapid review of a corpus of documents.

The present invention is equally useful within an electronic mail context where the user can view a summary of the electronic mail received and can then determine whether more or less of the contents of the entire electronic mail message(s) is desired.

**[0053]** Another useful application of the present invention is within the user interface of a modern computer system, such as the Apple Macintosh Finder, where stored documents (either locally stored, e.g., on a hard disk drive of the computer, or remotely stored, e.g., across a network or even across the internet) can be displayed by name, application type, date created, etc. When using such an interface, a user is oftentimes faced with a window displaying a long list of such stored documents without much hint as to what the documents actually contain. While documents or files are often given a particular name in order to provide a hint of their content or subject matter, the user is still often left wondering what a particular document or documents contain. As such, using the summarization engine of the present invention, the system could provide a “show top sentence” option. This option would display to the user the one sentence of a document which is most indicative of the contents of that document.

**[0054]** Such display could take the form of a portion of the display line or listing of documents in a computer system user interface as in a Finder folder window of the Macintosh computer system as is shown in **FIG. 6** wherein the amount of the top sentence displayed is limited by the amount of window display space allotted to this field. Such display could also take the form of being displayed in a comments field of an informational window about the document in a computer system user interface as is shown in **FIG. 7**. Such display could also take the form of being an expanded display in a display line or listing of documents when the user positioned a pointer over the document name or icon, when in a particular expanded display mode, or when

depressing a particular keyboard key and/or mouse button combination, as is shown in **FIG. 8**. Still further, such display could also take the form of an open dialog box where, instead of displaying a thumbnail miniature image of a graphic image document or merely the first sentence of a textual document, a summary comprised of a top sentence or sentences could be displayed, as is shown in **FIG. 9**.

**[0055]** An additional feature of the user interface document summary mechanism is the option, as in the more general document summary invention described above, for the user to control whether more or less of the document summary is to be displayed. In other words, while the default setting of a graphical user interface which displayed the “show top sentence” option might typically be to show only the one top sentence, the user could have the option of displaying a greater number of representative sentences from the summarized document. Such additional sentences might simply wrap onto the next line of the display or, instead, might only be displayed when the user positioned a pointer over the document name or icon when in a particular mode (e.g., similarly to the standard Macintosh Finder Balloon Help feature) or when depressing a particular keyboard key and/or mouse button combination. A large variety of display options is thus possible with the approach of the present invention depending upon such factors as display size and resolution, user preferences, and system capabilities.

**[0056]** In the foregoing specification, the present invention has been described with reference to a specific exemplary embodiment and alternative embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specifications and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

---

#### Appendix A

---

```
// -----
// find_next_sentence
// On return, start_of_sent will be > out if first chars encountered
// are whitespace.
//
// Normally returns length of sentence, starting from returned value
// of start_of_sent
// If it returns 0, then it ran out of buffer before finding
// a sentence. The caller will typically copy the remaining
// text to the beginning of a buffer, fill up the buffer,
// and then call this again. The case where a complete
// sentence does not fit in the buffer should be checked
// by the caller.
//
// Can't handle "see J.P. there?" or "call A. Morgan's"
// Handles Mr., Mrs., Ms., Dr., and i.e.
// -----
int find_next_sentence(char* buf, uint32 length, char** start_of_sent,
    bool *ran_out_of_buffer, bool *first_in_paragraph,
    bool remove_returns)
{
    *first_in_paragraph = False;
    // chew up leading whitespace
    char* last_loc_of_buffer = buf + length - 1;
    // identify if this is the start of a paragraph
    bool last_was_return = False;
    while (isspace(*buf) && (buf<=last_loc_of_buffer))
```

-continued

## Appendix A

---

```

{
    switch(*buf)
    {
        case '\r':
        case '\n':
            if(last_was_return) // return followed by return
                *first_in_paragraph = True;
            else
                last_was_return = True;
            creak;
        case '\t':
            if (last_was_return) // return followed by tab
                *first_in_paragraph = True;
            else
                last_was_return = False; // something came after the preceding
                // return other than a return or tab
                // break;
        case ' ':
            if (last_was_return && isspace(*(buf+1))) // return followed by
                // more than one
                // white space
                *first_in_paragraph = True;
            else
                last_was_return = False; // something came after the preceding
                // return other than a return or tab
            break;
        default:
            creak;
    }
    ++buf;
}
*start_of_sent = buf;
*ran_out_of_buffer = True;
if(buf>last_loc_of_buffer)
{
    *start_of_sent = 0;
    return 0; // note that past this point, we'll return *sum* length,
} // even if we hit end of the buffer before concluding a sent.
//Now we start looking for the end of the sentence.
*start_of_sent = buf;
bool conclusive_sentence = False;
bool abbrev = False;
char* lookahead;
do // we're going to repeat a big loop until we find a sentence break or
    // run out of characters in the buffer.
{
    switch(*buf) // Consider the current character in the buffer.
    case '"':
        if(*buf-1)=='.' // handle {Suzanne said "I love you." }
            // If it's a quotation mark preceded by a period,
            // we found a sentence break.
            conclusive_sentence = True;
            break;
        case '.':
            lookahead = buf+1; // If it's a period, consider next character.
            if(*lookahead == '.') // handle ellipses
                // If part of an ellipsis ( . . . ), consider
                // the character after the last period.
            {
                while( (*lookahead == '.') &&
                    lookahead <= last_loc_of_buffer)
                    ++lookahead;
            }
            if(lookahead > last_loc_of_buffer) // no more characters
            { buf = lookahead; break; }
            // rule out some abbreviations by checking for
            // space followed by capital letter
            bool was_space_after_period = False;
            while( isspace(*lookahead) && // skip white space
                // Was there a space after the
                // period? If so, it might be a
                // sentence break.
                lookahead <= last_loc_of_buffer)
            {
                ++lookahead;
                was_space_after_period = True;
            }
            if(lookahead > last_loc_of_buffer)
            { buf = lookahead; break; }
            if (!was_space_after_period)

```

---

-continued

## Appendix A

```

break;
// things a sentence can start with here
// If we have a quote, bullet, or dash after the
// space, we'll treat this as a sentence break.
if( *lookahead == '"' || *lookahead == '•' || *lookahead == '-' )
    ( conclusive_sentence = True; break; }
else if (!isupper(*lookahead)) break;
// If lowercase letter after period, it's not a sentence break.
// otherwise, check if it was just an abbreviation
// now we check for 'Mr.', 'Mrs.' etc.
// currently handles {Dr. Mr. Mrs. Ms. i.e.}
if(buf - *start_of_sent >= 2)
{
    switch(*(buf-1))
    {
        case 'r':    if(*(buf-2) == 'M' || (*(buf-2) == 'D')) // Dr. || Mr.
                        abbrev = True;
                        break;
        case 's':    if (*(buf-2) == 'M') // Ms.
                        abbrev = True;
                        break;
        case 'e':    if (*(buf-2) == '.' && *(buf-3) == 'i') // i.e.
                        abbrev = True;
                        break;
    }
}
if(buf - *start_of_sent >= 3    &&
    *(buf-1) == 's'    &&
    *(buf-2) == 'r'    &&
    *(buf-3) == 'M')
    abbrev = True;
// special case: if a period is immediately followed by a double quote
// count the quote as part of the sentence.
//if(!abbrev && *(buf+1) == '"')
// ++buff;
conclusive_sentence = abbrev; // if we get here its the
// simple case of end of
// sentence.
break; // that is "hello there. Go away now."

// catch & separate list items here (expensive)
// back to our initial character. If it wasn't a quote or period, what was it?
case '\r':
// This section is trying to separate lists of items (e.g., bullets) that may not
// use punctuation to separate the items.
case '\n':    if(remove_returns)
                *buf = ' '; // replace the return with a space
                lookahead = buf+1;
                while((*lookahead == ' ') && // skip space that might be
                    // between two returns
                    lookahead <= last_loc_of_buffer)
                    ++lookahead;
                if(lookahead > last_loc_of_buffer)
                { buf = lookahead; break; }
                // detect list items (lacking sentence punctuation clues)
                // If the newline followed by another, or a tab, or 3 or more
                // spaces, it's a sentence break.
                if(
                    *lookahead == '\n' || *lookahead == '\r' //two returns,
                    || *lookahead == '\t' // return followed by a tab
                    // --> paragraph delimiter
                    || (*lookahead == ' ' //return followed by 3 or more spaces
                        && *(lookahead+1) == ' ' && *(lookahead+2) == ' ')
                {
                    conclusive_sentence = True;
                    break;
                }
            }
            while( isspace(*lookahead) && // skip white space
                lookahead <= last_loc_of_buffer)
                ++lookahead;
            if(lookahead > last_loc_of_buffer)
            { buf = lookahead; break; }
            // Ditto if followed by a bullet or two hpyhens.
            if(*lookahead == '•' || (*lookahead == '-' && *(lookahead+1) == '-'))
            { conclusive_sentence = True; break; }
            break;

```



-continued

## Appendix A

---

```

// Back to our initial character. If a question mark or exclamation point,
// it's a break.
case '?':
case '!':
    conclusive_sentence = True;
    // if a period, '!', or '?' is immediately followed by a double quote
    // count the quote as part of the sentence.
    if(*(buf+1)=='"')
        buf++;
    break;
default: break;
}
buf++;
}
while ( !conclusive_sentence && (buf<=last_loc_of_buffer));
*ran_out_of_buffer = 'conclusive_sentence';
// return the length you conclusive_sentence, even if we ran out of buffer
before
    // determining conclusively whether it's a sent or not.
    // ran_out_of_buffer gives that indicator.
return buf - *start_of_sent;
}

```

---

We claim:

1. A document summary display system comprising:
  - a document containing one or more separate sentences;
  - a relevance ranking means for ranking the relevance of the one or more separate sentences of the document;
  - a display means for displaying a summary of the document within a computer system user interface listing containing a displayed reference to the document wherein the summary of the document is based upon the relevance ranking of the one or more sentences.
2. The system of claim 1 wherein the relevance ranking means ranks the relevance of the one or more separate sentences of the document to the document as a whole.
3. The system of claim 2 wherein the displayed document summary is comprised of one sentence of the one or more separate sentences of the document deemed most relevant according to the relevance ranking means.
4. The system of claim 2 wherein the displayed document summary is comprised of multiple sentences of the one or more separate sentences deemed most relevant according to the relevance ranking means.
5. The system of claim 4 wherein the number of multiple sentences of the one or more separate sentences in the displayed document summary is user controllable.
6. A document summary display user interface in a computer system comprising:
  - a document containing one or more separate textual portions;
  - a ranking means for ranking the one or more separate textual portions of the document;
  - a display means for displaying a summary of the document within a document information window of the computer system user interface wherein the summary of the document is based upon the ranking of the one or more separate textual portions of the document.
7. The system of claim 6 wherein the ranking means ranks the relevance of the one or more separate textual portions of the document to the document as a whole.

8. The system of claim 7 wherein the displayed document summary is comprised of one textual portion of the one or more separate textual portions of the document deemed most relevant according to the ranking means.

9. The system of claim 7 wherein the displayed document summary is comprised of more than one of the one or more separate textual portions of the document deemed most relevant according to the ranking means.

10. The system of claim 19 wherein the number of the one or more separate textual portions of the document in the displayed document summary is user controllable.

11. In a computer system user interface display comprising a listing of one or more documents, each document comprised of one or more sentences, the user interface display listing of one or documents also displaying one or more associated document parameters for each of the one or more documents, a document summary display comprising:

- a relevance ranking means for ranking the relevance of the one or more sentences of each of the one or more documents to the document comprising the one or more sentences;

- a display means for displaying, adjacent the one or more associated document parameters in the user interface listing of the one or more documents, a summary of each of the one or more documents wherein the summary of each of the one or more documents is based upon the ranked relevance of the one or more sentences of each of the one or more documents.

12. In a computer file system file directory listing displayed on a computer display, a document summary display comprising:

- a document containing one or more separate textual portions;

- a ranking means for ranking the one or more separate textual portions of the document;

- a display means for displaying a listing of the document within the computer file system file directory and also for displaying a summary of the document within the

computer file system file directory listing containing the document listing wherein the summary of the document is based upon the ranking of the one or more separate textual portions of the document.

**13.** The document summary display of claim 12 wherein the ranking by the ranking means is based on relevance of the separate textual portions to the document as a whole.

**14.** The document summary display of claim 13 wherein the separate textual portions of the document are sentences contained in the document.

**15.** The document summary display of claim 14 wherein the summary of the document contains a portion of the highest ranking sentence of the document.

**16.** In a computer file system file directory listing displayed on a computer display, a method of displaying a summary of a document comprising one or more sentences, the method comprising:

displaying in the computer file system file directory a listing referencing the document;

ranking the relevance of the one or more sentences of the document to the document as a whole:

displaying in the computer file system file directory the highest ranking one or more sentences of the document.

**17.** An electronic user interface document summary method comprising:

displaying in the electronic user interface a reference to at least one document, the document comprised of at least one textual portion;

ranking the at least one textual portion of the at least one document;

displaying a summary of the at least one document in the electronic user interface, the summary comprised of the highest ranking textual portion of the at least one textual portion of the at least one document.

**18.** The electronic user interface document summary method of claim 17 wherein the ranking is by relevance of the at least one textual portion of the document to the at least one document as a whole.

**19.** The electronic user interface document summary method of claim 18 wherein the at least one textual portion of the at least one document is a sentence of the at least one document.

**20.** The electronic user interface document summary method of claim 19 wherein the reference to the at least one document in the electronic user interface is in a window of the electronic user interface listing the at least one document.

**21.** The electronic user interface document summary method of claim 19 wherein the reference to the at least one document in the electronic user interface is in a window of the electronic user interface displaying information about the at least one document.

\* \* \* \* \*